

1st HPC User Meeting

Rainer Krug, Albert Groenwold, Evan Lezar

11 November 2009

The first user meeting of the Stellenbosch University HPC cluster was held on 11 November 2009. The chair for the meeting was Rainer Krug with Evan Lezar and Albert Groenwold providing assistance. This document serves to summarise most of the important points raised or addressed during the meeting.

1 Attendance

The meeting was attended by 14 users from various faculties and departments. The list of attendance is as follows:

Name	Department
Beatrix Coetzee	Genetics
Heine de Jager	IT
Andrew de Wett	M&M
Jaco Geldenhuys	Computer Science
Albert Groenwold	M&M
Dan Jacobson	Viticulture
Rainer Krug	CiB (Bot-Zoo)
Evan Lezar	E&E
Andries Nieuwoudt	IT (+SC)
Kenneth Oberlander	Consent
Anne Ropiquet	Bot-Zoo
Salomon van Huyssteen	M&M
Gerhard Venter	M&M (+SC)
Gideon Wiid	E&E

2 Report from Steering Committee Meeting (Gerhard Venter)

A quick run-through of the minutes of the last steering committee meeting was given, with a number of points leading to further discussion. These points will be covered in subsequent sections where applicable.

3 Resource use of Rhasatsha (Evan Lezar)

3.1 Queueing

The question of queueing led to quite a bit of discussion. It was felt by all present that the tiered model presented by the steering committee is a fair one and that some preference should be given to departments that have contributed in the way of funds or hardware. The exact waiting of these priorities may still need some refinement, but requires that the current values and usage be monitored and user feedback obtained if one feels that the queueing system is unfair.

The fact that a general facility such as this is home to a wide range of users, with just as wide a range of requirements makes the refinement of the queueing process a none trivial one. Users are reminded that at any such facility there will be a queue with many such facilities requiring users to apply for CPU time weeks in advance. That said, if users feel that they are being unfairly treated, they are requested to contact a member of the steering committee as a first step and to refrain from resorting to using the public mailing lists to resolve such matters.

A possible adjustment to the queueing system that was suggested is to make provision for a runtime estimation on the side of the user as part of the submission system. This has the advantage that the queueing system may be better able to manage the available resources. It does however complicate matters for the users and for some types of experiments it is difficult if not impossible to perform such estimates.

It was decided that a monthly summary of usage statistics for all users be generated and posted on the wiki and/or the mailing list. This will give a better idea of who is using the cluster and whether the implemented queueing strategies are affecting users adversely.

Another point that was raised was the availability of systems campus-wide that are idle most of the time, but since these are production-level systems, database servers for example, there are times when their computational power is required for a specific task. This also falls into the category of cluster expansion, with the question being whether there is space in the current models to support guaranteed CPU time for certain tasks with hard deadlines.

3.2 Disk usage

Disk space is the other shared usage that has to be fairly distributed. Although it is not favourable to set up a quota system due to variable demands by different user applications, it is felt that something should be done to prevent users from leaving data on the cluster for extended periods of time.

One suggestion is that a small limit is placed on an account for the minimum yearly subscription fee (say 5GB or 10GB), with a small additional fee allowing users "unlimited" storage space. The idea behind this is that the serious users will be willing to spend the extra money and effort to ensure that their accounts are fully functional, whereas users who use the system infrequently will not be able to archive large amounts of data. It was decided to try this approach from next year, although the lower limit as well as the additional cost still needs to be decided.

The question of donation of personal hard drives for storage space was also raised. One problem with

this is the fact that the storage subsystem makes use of SAS drives in a dedicated enclosure, which is currently at capacity. Expansion of storage capacity thus requires the purchase of additional hardware over and above the physical disks increasing the initial cost per gigabyte. It was mentioned that a tiered storage scheme involving slower cheaper disks could be possible, but this would come at an increased cost in terms of administration and a possible degradation in performance for some users.

A problem with getting large amounts of data off the cluster was also raised in that the campus network is limited to 100Mbps. This translates to a transfer time of about 2 minutes per gigabyte. One suggestion is to have a terminal at IT available for the transfer of data. If this terminal is connected directly to the gigabit switching hardware and supplies a number of USB ports or even eSATA ports this would greatly reduce the time required for large datasets. That said, a user could easily start a transfer overnight and copy almost 300GB of data in an 8 hour period.

As is the case in queue usage, it is recommended that weekly or monthly reports of disk usage are sent to the mailing list or posted on the wiki. This will serve to remind users to clear out their home folders whenever possible.

4 Representation of users on steering committee (Evan Lezar)

The question was raised as to the users felt that their representation on the steering committee was adequate, or alternatively put, that the steering committee is catering to their needs. After some discussion it was decided that no additional user representation on the steering committee is required. It is requested, however, that a member of the steering committee be chosen as the user liaison and that this person serve as the person of contact for issues raised by the users.

Regarding the steering committee meetings, it is requested that the users be given at least 2 weeks notice of a meeting and that the tentative agenda also be made available ahead of time. This will afford users the time to formulate any comments or questions that they may have for the steering committee and also allow for organising a user meeting if required. The minutes of the meetings must then also be made available to the users on the HPC wiki.

5 Funding ideas for cluster expansion (Albert Groenwold)

One possible source of funds for the cluster is the direct donation from individual users or research groups using the cluster. Just prior to the meeting it was announced that a researcher had contributed a sum from their own research budget to the central pool which is managed by the steering committee. Regarding the question of procedure for this, it is recommended that a researcher or group wishing to make a contribution contact the steering committee to do so.

Regarding the types of donations, these can be in the form of funds - either to be used to buy specific hardware - or to be placed in a central pool and used for future upgrades and maintenance. Another option is for users to donate hardware directly, especially if special-purpose hardware is required. If this path is followed, it is important to ensure that the hardware is compatible with the existing installation. As discussed in Section 3.1, it may be that some of this hardware is donated on the grounds that certain

users have priority in its usage - especially for mission critical applications. Some where however of the opinion that hardware donated to the cluster in this way would be managed along the lines of the same fair use policy used thus far and that no preferential treatment (apart from increased queue priorities as with cash contributions) should be given. Users with questions regarding this are welcome to contact the steering committee, or Prof Venter specifically.

It was also put forward that M&M alone has allocated a large sum (more than what the original cluster costed) for future cluster expansion. However, if no other departments or research groups make a contribution to the second phase of the project, the logical step is for them to start their own facility - to the detriment of all current users and the HPC community at the university as a whole. It is therefore imperative that other parties do start committing to contributions to continue the current shared model.

One of the suggestions for increasing the funding is to simply increase the yearly subscription fee. Although this may help for maintenance of the cluster, it is not likely that this will assist with increasing capacity to a great extent. Furthermore, if the fee is made too high, the likelihood that the usage of the cluster will continue to increase is lowered which may not be a favourable outcome. On the contrary, the point of the central facility is to get more faculties involved.

It was felt that ideally the facility would be funded to a large extent from the university itself, although it is unclear how best to approach this. One of the options is to propose a model where the HPC facility gets a small percentage of any revenue streams from subcommittee A or B. It is suggested that this model be investigated and the figures and a proposal be presented to Prof van Zyl. Another option is to approach the CHPC and see if providing funding for a university facility falls into their vision for research and human capital development. The question of the use central funds for equipment at a faculty or university level was also raised.

In all the scenarios discussed an important factor in selling the idea of HPC to the top level management at the university is to show that the facility is an invaluable research tool. It is thus suggested that the compilation of a list of publications that used the facility be made a top priority. This is to include a list of publications (e.g.. journal articles, conference proceedings) as well as theses (masters and doctoral dissertations). Furthermore, the works should be categorised into sections depending on whether the HPC facility was essential to the research or whether the work only benefited greatly. A third category could include work that may have future HPC applications if the facility had the right infrastructure. A list of these publications could also be kept on the wiki. With this in mind, it is suggested that a standard citation for acknowledging the HPC facility be made available on the wiki for use in works that are made possible by the facility.

6 Additional points from Steering Committee Meeting (Albert Groenwold)

One of the points that was raised at the steering committee meeting and clarified at the user meeting was the question of internet access, and specifically access to facilities such as the CHPC. At present the total bandwidth of the university is at 50Mbps, with an upgrade to 1Gbps taking place by the beginning of December which includes 600Mbps international bandwidth. The remaining 400Mbps would be available for access to the SANREN network, for example. In terms of the new capacity, much of the infrastructure is already in place and in the process of being commissioned with the final step being the provision of

the fibre interlink.

In terms of pricing and especially making provision for free traffic to the CHPC and other institutions, the model followed will depend on the network management and internet model used at the target sites. In the past for example, free links between institutions were used to simply proxy internet traffic and bypass the existing university billing facilities. The infrastructure being installed does however make provision for granting access to certain destinations (such as the CHPC) using a different funding model from general internet access.

7 Regular user meetings (Rainer Krug)

It was decided that user meetings will be held every quarter (3 months) and would in future include some short user presentations showcasing work or methods applicable to the HPC. Thus there will be no separate mini-symposia at present. It is also recommended that the user meetings coincide with the steering committee meetings so as to facilitate feedback between the two. It is requested that the steering committee give the users at least 2 weeks notice of a steering committee meeting so that a user meeting can be arranged and questions, concerns and recommendations raised. There is also the possibility of a short feedback session following the steering committee meeting if pressing issues need to be addressed.

In terms of the provision of snacks and refreshments at the meeting, it was decided that the current model of the users themselves contributing to this be followed. It is however recommended that in future a single user be responsible for the collection of the pizza and beer (for example) so as to minimize the administration involved. The cost will then be divided between the users willing to contribute.

The question of an informal gathering such as an end of year braai was also raised. It was decided that the issue will be taken further on the mailing list and wiki.

8 Who is doing what on the cluster? (Rainer Krug)

In such a community there is often a question as to what everybody is doing. A quick pass was made around the room and brief introductions given. It is requested that all users try to maintain a page on the HPC wiki where they list their affiliation as well as the kind of work that they are interested in. An example page will be sent to the mailing list to help new users with the process.

9 GPGPU computing at Stellenbosch (Evan Lezar)

An interest group for general purpose GPU computing at Stellenbosch was introduced. The aim behind the group is to bring together researchers from various fields all interested in GPU computing with hope of fostering collaborations as well as sharing of resources, such as testing platforms, and experience.

To facilitate communication, a mailing list **maties.gpgpu** was started and a wiki page (GPGPU) on the HPC wiki created.

10 Diverse (Rainer Krug)

No points were raised under this heading.

11 Closing thoughts

It is important that we as users be proactive in the use and management of the facility. The community that is developing around its use is an active and energetic one and we would like to keep it this way. Thus make an effort to familiarise yourselves with the wiki and to update your pages and the list of publications. Also keep an open mind for collaborations as it is really is a golden opportunity to do some great interdisciplinary work with HPC as the common ground.